arXiv:1610.05747v1 [stat.ME] 18 Oct 2016

Modeling Heterogeneous Peer Assortment Effects using Finite Mixture Exponential Random Graph Models

Teague Henry, Kathleen Gates, Mitchell Prinstein
University of North Carolina at Chapel Hill
Douglas Steinley
University of Missouri

Submitted to Psychometrika

Abstract

This article develops a class of models called Sender/Receiver Finite Mixture Exponential Random Graph Models (SRFM-ERGMs) that enables inference on networks. This class of models extends the existing Exponential Random Graph Modeling framework to allow analysts to model unobserved heterogeneity in the effects of nodal covariates and network features. An empirical example regarding substance use among adolescents is presented. Simulations across a variety of conditions are used to evaluate the performance of this technique. We conclude that that unobserved heterogeneity in effects of nodal covariates can be a major cause of mis-fit in network models, and the SRFM-ERGM approach can alleviate this misfit. Implications for the analysis of social networks in psychological science are discussed.

1 Introduction

Psychologists and social scientists at large increasingly utilize methods to analyze social networks. These methods typically attempt to quantify how individuals behaviors or traits relate to specific peer selection processes, depicting how a given individuals behavior or trait changes the network structure over time. The cross-sectional pattern of peer nominations conditioned upon individual traits and covariates is called "peer assortment". The relationships between covariates and peer assortment present complex methodological problems. One such problem is that heterogeneity in the peer selection and influence processes likely exist. Specifically, some individuals may have different patterns of relations among behaviors, traits, and peer selection. The present paper introduces a novel method that utilizes finite mixture modeling within an exponential random graph modeling framework to allow for heterogeneity in the patterns among covariates and peer selection or influence.

In recent years, new methodologies have been developed to analyze peer selection processes (e.g., SIENA; [1]; STERGMs; [2]). Models that investigate cross-sectional peer assortment are similar to models that assess peer selection processes, specifically the exponential random graph modeling framework (ERGM; [3]; [4]). What underlies all of these network analytic methods is their ability to assess the effect of individual level traits and covariates, such as personality characteristics or observed behaviors, on the network structure. In directed networks (e.g., peer friendship networks where individuals nominate peers as friends, but there is no requirement for bidirectional nomination), researchers can assess the effects of covariates on nominating a peer (sending) or being nominated (receiving). These effects assess the impact of a behavior, such as substance use, on the tendency for individuals to send or receive edges (i.e., nominate or be nominated). Identifying covariate effects of sender and receiver components prove useful when investigating the impact of individuals behaviors on network structure.

Parallel to network methods recent flourishing, methods that seek to model and capture individual (or ideographic) differences are also resurging and may benefit from a network approach. Modeling ideographic differences can be done a variety of different ways. For example, an ideographic analysis of relations of effects [5] allows researchers to examine a single individual's pattern of behavior, while other methods test the aggregate relations between outcomes as moderated by some individual trait or behavior [6]. Of the existing individual differences methods, moderation by an observed variable is by far the most commonly used [6]. Indeed, neglecting to include moderation effects in non-network data may result in biased estimates [7] and spurious random effects [8]. It follows that failure to account for heterogeneity in network analysis could have potentially severe consequences, especially because certain parameter combinations can cause the estimation to converge on a degenerate solution [9]. This risk of degeneracy could potentially increase when estimating a mis-specified model.

Given burgeoning interests in networks and individual differences, a growing number of researchers are seeking to examine the role of moderators and individual differences in networks (for review see, [10]), with the majority of attention going to peer influence and little work exploring potential moderators of peer assortment (with exception; see [11]). Despite the interest in the field, there currently are no methodological approaches within network science for moderation of these effects based on unobserved classes. There are several statistical methods for detecting latent com-

munities, such as stochastic block modeling [12], latent space modeling [13], a latent space model that allows for covariate effects [14], a mixture latent community approach [15], and more recently, the ERGM approach [16] that allows for differences in the effects of covariates within a community, and a stochastic block modeling approach that also allows for different effects of covariates within a community [17]. However, all of these approaches focus on assessing latent community structure and do not focus on assessing individual differences in the effects of covariates based on unobserved subgroups. Clusterwise p^* modeling [18] also focuses on determining latent community structure, but has the feature of using an ERGM model to inform the assignment of nodes to the latent communities. The method developed in this manuscript bears resemblance to the clusterwise p^* model, however does not require that the latent class structure result in interconnected latent communities. In this way it can be looked at as a generalization of the work in [18].

The present manuscript aims to extend the definition of mixtures on networks, and model mixtures of individuals who differ in their effect of observed variables on their position in the network. Latent class models and mixture models, which are also commonly used in psychological research to assess individual differences based on unobserved data, may provide a solution to this methodological hindrance. Specifically, finite mixture linear regression [19] and finite mixture general linear regression [20] can be used to partition a sample into groups, in which each group has different relations among observed variables. Following this line of research, this manuscript develops a sender/receiver finite mixture modeling approach for ERGMs (SRFM-ERGM), focusing on cross-sectional peer assortment models. This work will help address the methodological conundrum facing researchers who wish to utilize a whole network analysis method to test the moderation of peer assortment by individual differences. We first present sufficient details of ERGMs as well as the methodological innovation. Next, we present a set of simulations that demonstrate the SRFM-ERGMs ability to recover the mixture components, as well as the consequences of failing to model the individual differences. Finally, an empirical example examining individual differences in the effect of alcohol use on the network structure in a middle-school will be presented.

2 Network Terminology and Model Description

As social networks typically take the form of binary (in that a social relationship is either present or not), and directed (where a social relationship is not necessarily reciprocated), the description of network terminology will be specific to binary directed networks. A network consists of a set of *nodes* and a set of *edges* that connect the nodes. Nodes define the members of a network and, in the case of social networks, usually represent people. Edges define the relationship between nodes in the network. In the case of social networks, this is a social relation, such as friendship. In a directed network, relations do not need to be reciprocated, as individual i might nominate individual j, while individual j might not nominate individual i. An adjacency matrix, $\mathbf{A}$, is a $N \times N$ matrix where where N is the number of individuals in that comprise the network and A_{ij} is 1 if individual i has nominated individual j and 0 otherwise.

In addition to the presence or absence of edges, often times *node-level attributes* are also included in the analysis. In this manuscript, the $N \times P$ matrix of node level variables will be denoted using $\mathbf{X}$.

3 Exponential Random Graph Models

ERGMs, also referred to as p^* models, [4] are a family of models used to analyze network structure both contingent on endogenous predictors as well as on covariates of the individual nodes. ERGMs model the entire network as a sample from a population of networks with a set of sufficient statistics associated with the population. These sufficient statistics can include the number of edges in the network, the number of edges between individuals with dissimilar values on one specific covariate and/or the number of closed triangles in the network, to name a few examples. ERGMs were originally developed to analyze *dyad independent* networks, where the probability of an edge between two nodes was only dependent on the characteristics of those two nodes [21]. They were extended by [3] to include transitivity effects, where the probability of an edge existing depends on the existence of neighboring edges. The modeling of transitivity proves to be difficult, and only in recent years have successful modeling techniques been developed [22].

The SRFM-ERGM developed in this manuscript is a class of models posits that every node has a latent component assignment, and this component assignment moderates the effect of specific nodal covariates on the probability of edges directed towards the node (A Receiver Finite Mixture ERGM), or edges directed away from the node (A Sender Finite Mixture ERGM), but not both. This division often aligns with hypotheses made by researchers who are interested in associations between nodal attributes and nominations. The choice to restrict the model to either assess the effect of either sender latent class or receiver latent class was made to render the model tractable to estimation via a classification likelihood EM-type algorithm [23, 24, 25] for reasons elaborated below.

An ERGM models the observed network as a sample from a exponential family distribution with the following form. Let $\mathbf{A}$ be a random directed network of size N with sample space $\mathcal{A}$, and $\mathbf{X}$ a $N \times P$ matrix of fixed and known nodal covariates, and with $\mathbf{a}$ being a sample realization of $\mathbf{A}$, with a_{ij} being the sample realization of A_{ij}. Then the probability density function of $\mathbf{A}$ is

$$P(\mathbf{A} = \mathbf{a}|\mathbf{X}) = \frac{\exp[\langle \boldsymbol{\theta}, s(\mathbf{a}, \mathbf{X})\rangle]}{\psi(\boldsymbol{\theta})}. \tag{1}$$

Where $\boldsymbol{\theta}$ is a d length vector of natural parameters in log-odds metric, $s(\mathbf{a}, \mathbf{X})$ is a d length vector of sufficient statistics, $\langle \boldsymbol{\theta}, s(\mathbf{a}, \mathbf{X})\rangle$ is the inner product of the vector of natural parameters and the vector of sufficient statistics and $\psi(\boldsymbol{\theta})$ is a normalizing constant. The log-likelihood of the parameters [22] is then:

$$\ell(\theta) = \langle \boldsymbol{\theta}, s(\mathbf{a}, \mathbf{X})\rangle - \log(\psi(\boldsymbol{\theta})). \tag{2}$$

.

As was outlined above, the vector of sufficient statistics $s(\mathbf{a}, \mathbf{X})$ can contain a variety of terms. Sufficient statistics that will be used in the simulation studies and the empirical examples are described below. Broadly, these can be divided into two categories, sufficient statistics that are wholly based on the network structure, *structural statistics*, and statistics that are based on the node level attributes, *node level attribute statistics*.

3.1 Structural Statistics

Edges. The parameter estimate associated with the edges statistic in any ERGM acts as an intercept for the model. The sufficient statistic described by this term is a count of the number of edges in the network.

Mutuality. The sufficient statistic is the total number of reciprocated dyads (Holland & Leinhart, 1981). These are dyads where $a_{ij} = a_{ji} = 1$. The parameter associated with this statistic can be described as the tendency for edges to be reciprocated.

Geometrically Weighted Edgewise Shared Partners (GWESP). This sufficient statistic is used to account for transitivity in the network [26]. Transitivity is the phenomena where an edge is more likely to be present if the edge has shared partners. This is commonly referred to colloquially as "A friend of a friend is my friend." Transitivity is a common phenomena in real-world networks, and failure to model transitivity can result in bias in the estimates of other effects [27]. However, when a term that models transitivity is added, this induces a dependency structure in the network that prevents a direct maximization of the likelihood during estimation. This dependency structure also runs the risk of causing degeneracy in estimation, where the expected value of the network's distribution converges to a completely empty or completely full network. Details of why this occurs can be found in [9]. The geometrically weighted edgewise shared partners term works to prevent degeneracy by down-weighting large numbers of shared partners. A GWESP term has both an effect parameter $\theta^{(GWESP)}$ and a weight parameter τ. τ ranges from 0 to ∞, with 0 meaning that the effect of shared partners does not increase beyond having 1 shared partner, and ∞ meaning that the effect of additional shared partners is linearly increasing with sample size.

If τ is allowed to be freely estimated, then the ERGM becomes a member of a curved exponential family [28], which complicates estimation. In the simulations that follow, τ is considered fixed and known, a practice that is done in substantive analysis where the curved exponential family models fail to converge [26].

3.2 Nodal Attribute Terms

The inclusion of nodal level covariates is described in detail in [29]. These terms allow an analyst to model the effects of nodal level covariates on the probability of an edge being present for both an individual node level, as well as for dyads.

Attribute Matching. The parameter estimate associated with this sufficient statistic is used to evaluate the effect of nodes matching on a categorical variable (such as gender or ethnicity) on the probability of having an edge . The sufficient statistic is the number of edges between nodes that match on the categorical variable.

Attribute Outdegree The parameter estimate associated with this sufficient statistic is used to evaluate the interaction effect of a continuous variable and the number of edges coming out of nodes on the probability of having an edge. The node level sufficient statistic is the outdegree of a node multiplied by the attribute value for that node.

Attribute Indegree Similarly to the Attribute Outdegree term, the parameter estimate associated with this sufficient statistic reflects how a continuous variable differentially relates to the probability of an edge based on the number of edges the node receives. The node level sufficient statistic is

indegree of the node multiplied by the attribute value for that node.

Attribute Absolute Difference The parameter associated with this sufficient statistic evaluates the effect of similarity on a continuous variable between nodes on the probability of an edge. The edge-level sufficient statistic is the absolute difference in the values of the attribute for the sender and receiver of the edge multiplied by the presence or absence of that edge.

4 Estimation for ERGMs

Estimation of the parameter vector $\boldsymbol{\theta}$ for an ERGM model is complicated when a term that measures transitivity is included in the model. ERGMs, as originally proposed by [21], did not contain transitivity terms, and instead were considered dyad independence models. Dyad independence models are ones where the state of a dyad (what edges in a dyad are active), are conditionally independent of all other dyads given the characteristics of the nodes that make up the dyad in question. This simplifies the likelihood of a *dyad independent* ERGM to that of a multinomial regression for a directed network (the categories being the 4 possible states of a directed dyad), or a binary logistic regression for an undirected network, which in turn makes estimation of a dyad independence ERGM trivial.

When a term that accounts for transitivity, such as a GWESP term, is added to the model, the resulting dependency structure complicates estimation. Specifically, the likelihood changes due to a normalizing constant that is defined with a sum over all possible networks of a given size, which is an astronomically large number for any reasonable sized network.

A great deal of work has been done on developing approximate MLE estimators for ERGM models, and the current state of the art estimator is the MCMC sampler of [22]. This estimator uses MCMC to approximate the intractable normalizing constant. However due to the iterative nature of the [22] estimator, it is less than ideal to implement in an expectation maximization framework, which in of itself is iterative. Additionally, the MCMC sampler does not scale well to larger networks, limiting the practical size of a network to less than 300 nodes. An alternative to the ML estimator appears in the maximum pseudolikelihood approach of [3].

4.1 Maximum Pseudolikelihood Estimation of ERGMs

The *pseudo-likelihood* approximation to the log-likelihood as developed by [3] for an ERGM is then the likelihood for a binary logistic regression of the following form:

$$P(A = a|\boldsymbol{\theta}) \approx \prod_{(i,j)\in\mathbf{a}} \frac{(\exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle))^{a_{ij}}}{1 + \exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle)}. \tag{3}$$

Where $\mathbf{s}_{ij}$ is the change in the sufficient statistics when A_{ij} goes from 0 to 1. This expression is identical to a logistic regression model treating the $\mathbf{s}_{ij}$ as fixed and known covariates. Taking the log of the above transforms it to the pseudo-log-likelihood which can be simplified as shown:

$$\hat{\ell}(\boldsymbol{\theta}) = \log\left[\prod_{(i,j)\in\mathbf{a}} \frac{(\exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle))^{a_{ij}}}{1+\exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle)}\right] \tag{4}$$

$$= \log\left[\frac{\exp(\sum_{(i,j):a_{ij}=1}\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle)}{\prod_{(i,j)\in\mathbf{a}} 1+\exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle)}\right] \tag{5}$$

$$= \sum_{(i,j):a_{ij}=1} \langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle - \sum_{(i,j)\in\mathbf{a}} \log(1+\exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle). \tag{6}$$

Note that $\langle\boldsymbol{\theta}, s(\mathbf{a}, \mathbf{X})\rangle$ is equal to $\sum^{(k)} \theta^{(k)} s^{(k)}(\mathbf{a}, \mathbf{X})$, and that $s^{(k)}(\mathbf{a}, \mathbf{X})$ is equal to $\sum_{(i,j):a_{ij}=1} s_{ij}^{(k)}$. Therefore $\sum_{(i,j):a_{ij}=1}\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle = \langle\boldsymbol{\theta}, s(\mathbf{a}, \mathbf{X})\rangle$. This leads to the final expression of the pseudo-log-likelihood below:

$$\hat{\ell}(\boldsymbol{\theta}) = \sum_{(i,j)\in\mathbf{a}} \log(P(A_{ij}=1|\mathbf{A}_{ij}^c)) = \langle\boldsymbol{\theta}, s(\mathbf{a}, \mathbf{X})\rangle - \sum_{(i,j)\in\mathbf{a}} \log(1+\exp(\langle\boldsymbol{\theta}, \mathbf{s}_{ij}\rangle)). \tag{7}$$

Note that this pseudo-log-likelihood differs from the log-likelihood only in the normalizing constant. Furthermore, this likelihood is identical to that of a binary logistic regression, and can be maximized by established means.

There are several issues with using the MPLE for ERGMs that model transitivity. [30] provide evidence that the standard errors of the estimates in an MPL fit ERGM tend to be underestimated, and the whole model tends to overfit to the data. [27] show in a simulation study that the estimates of the transitivity parameter from an MPL fit ERGM are substantially less efficient than that of a ML fit ERGM. However, recently it has been proposed that the properties of the MPL estimator have not been thoroughly studied [9], and there is evidence for its asymptotic normality [31].

With the limitations of MPLE discussed, there are several advantages that suggest MPLE is a better choice than MLE for the current purposes. Using the MLE during the maximization step is computationally intensive, with running times an order of magnitude above the running time of MPLE. Furthermore, in light of the possibility of degeneration, repeated applications of the MLE while using non-optimal latent class solutions increases the probability that at least one iteration will be degenerate, which would cause estimation to fail. It is worthwhile to note that both of these issues were raised in [18], and motivated the use of the MPLE there as well. As such, in current manuscript, we use MPLE during the maximization step of an EM algorithm implementation. Given the known issues with MPLE, we use MLE to obtain a final set of parameter estimates, once the EM algorithm has converged.

5 Sender/Receiver Finite Mixture ERGMs

[20] developed a mixture model for generalized linear models. This approach relies on the EM type algorithm for estimation and provides an elegant solution to the problem. As the MPL for an exponential random graph model is identical to the likelihood for a binary logistic regression, [20] approach can be applied.

The joint log pseudolikelihood for a SRFM-ERGM is:

$$\hat{\ell}(\mathbf{Z}, \boldsymbol{\theta}) = \sum_i \sum_{q=1}^{Q} z_{iq} \log(\alpha_q) + \sum_i \sum_{j \neq i} \left[a_{ij} \left[\sum_{q=1}^{Q} z_{iq} \langle \boldsymbol{\theta}_q, s_{ij} \rangle \right] - \sum_{q=1}^{Q} z_{iq} \log(1 + \exp(\langle \boldsymbol{\theta}_q, \mathbf{s}_{ij} \rangle)) \right]. \quad (8)$$

Where Z is an $N \times Q$ matrix of indicators of class membership, where z_{iq} is 1 if node i is in class q and 0 otherwise, where $q = 1, 2, 3 \ldots, Q$. α_q is the marginal probability of a node being in class q. $\boldsymbol{\theta}_q$ is a vector of parameters for class q. a_{ij} is the value (0 or 1 for a binary network) of the edge from node i to node j. s_{ij} is the vector of changes in the sufficient statistics of the network if edge a_{ij} went from 0 to 1, given the rest of the network. Any parameter θ may be held equal across all classes, and thus rendered homogeneous.

The choice to restrict these models to mixture components defined by either sender or receiver nodes is due to the dependency structure noted by [15] in their development of mixture models for Bernoulli networks. Should the mixture component be defined by both incoming and outgoing ties then the vector θ_q would actually be θ_{ql} where l is the mixture component label of the receiving node. As such, each edge in the network would have a likelihood that is dependent on the mixture component of both the sending and receiving node. In turn, as was noted by [15], the expected value of a mixture component label for individual i becomes dependent on the component label of every other node in the network. This requires that the expected value of component label be estimated simultaneously for every node, leading to a sum over every possible combination of component labels in the denominator, which very quickly becomes computationally infeasible. The choice to restrict the models to either sender or receiver leads to a separable expected value statement for each node.

5.1 Expectation Step

The general expectation for a latent class general linear model (where the probability densities are exponential family) is in the following form ([20], Eq 7):

$$E[z_{iq}|\boldsymbol{\theta}, \mathbf{y}_i] = \frac{\alpha_q \prod_{k=1}^{K} f_{ik|q}(\mathbf{y}_{ik}|\theta_{\mathbf{q}})}{\sum_{q=1}^{Q} \alpha_q \prod_{k=1}^{K} f_{ik|q}(\mathbf{y}_{ik}|\theta_{\mathbf{q}})}. \quad (9)$$

Where $\mathbf{y}_i$ is a vector of observations (each of which can come from a different arbitrary probability density) for case i and $f_{ik|q}(\cdot)$ is the probability density function for the kth variable of the ith case given the ith case is in class q.

Here we can express the expectation for the sender latent class exponential random graph model in precisely that form of (9):

$$E[z_{iq}|\boldsymbol{\theta}, \mathbf{a}] = \frac{\alpha_q \exp(\sum_{j \neq i} [a_{ij} [\langle \boldsymbol{\theta}_q, s_{ij} \rangle] - \log(1 + \exp(\langle \boldsymbol{\theta}_q, s_{ij} \rangle))])}{\sum_{q=1}^{Q} \alpha_q \exp(\sum_{j \neq i} [a_{ij} [\langle \boldsymbol{\theta}_q, s_{ij} \rangle] - \log(1 + \exp(\langle \boldsymbol{\theta}_q, s_{ij} \rangle))])}. \quad (10)$$

Note that in the above likelihood, edges that share the same sender node all are assigned the same latent class. Intuitively, this is due to the assignment of the latent class being sender

based rather than edge based. This maps onto the Wedel and DeSarbo approach, as they classify observation vectors. In this case the observation vector is the vector of all possible edges a node could and did send.

With the expectation derived, we diverge from the approach laid out in [20]. In [20]'s approach, the expectations produce probabilities of class membership, and these probabilities are then used in the maximization step. This approach is known as mixture EM, which is the EM algorithm laid out originally by Dempster, Rubin and Laird [25]. In this manuscript, we take a classification likelihood approach. The classification likelihood approach (for examples see: [32, 33, 34]), assigns each case the class label that has the highest probability. This hard assignment of classes implies that although the likelihood of the data will be a mixture across the classes, each case is not a mixture across the latent classes. Classification likelihood can be fit using an EM-type algorithm that uses the class label with the highest probability instead of using the expected value of class membership (a mixture approach).

This hard assignment has advantages and disadvantages. Classification likelihood algorithms consistently find more informative latent classes than mixture algorithms in the sense that the Kullback-Liebler divergence [35] of the mixture components is maximized; however it does not maximize the likelihood of the data as the mixture EM algorithm does [23]. In terms of disadvantages, classification likelihood does not recover ill-separated mixture components well, nor does it recover very unbalanced group sizes [36, 37]. This failure to recover ill-separated groups comes from the maximization of the KL divergence. If the groups are ill-separated, classification likelihood tends to model the data using a single distribution in which the majority of cases are placed. This can be an advantage if the presence of latent classes is being tested; however, if the presence of latent classes is assumed, this approach will not detect them if they are ill-separated. This makes classification likelihood more conservative in its detection of latent classes, whereas mixture likelihood approaches could result in the over extraction of the latent classes [38]. Additionally classification likelihood does not exhibit optimal large sample properties in that it tends to be asympototically biased [39], though this bias appears to be lessened when allowing for different cluster sizes [36].

In this manuscript, classification likelihood is used for both its computational benefits (in that the likelihood easily separates on a case by case basis), as well as its property of maximum informativeness as measured by KL divergence. This choice is in part based on the reasoning presented in Steinley and Brusco [40], in that a mixture-likelihood approach is better for density estimation, while a classification likelihood approach is better for the clustering of observations. The finding of more informative clusters, as well as classification likelihoods tendency to be more conservative in its clustering are two major advantages of this approach. By using classification likelihood, the Sender/Receiver Latent Class ERGMs are explicitly attempting to classify individuals into well separated clusters, rather than approximate a complex distribution using mixtures of simpler distributions.

5.2 Maximization Step

With the latent class labels assigned, the conditional likelihood is:

$$\hat{\ell}(\boldsymbol{\theta}|\mathbf{Z},\mathbf{a}) = \sum_{(i,j)\in\mathbf{a}} \left[a_{ij} \left[\sum_q Z_{iq} \langle \boldsymbol{\theta}_q, \mathbf{s}_{ij} \rangle \right] - \sum_q Z_{iq} \log(1 + \exp(\langle \boldsymbol{\theta}_q, \mathbf{s}_{ij} \rangle)) \right], \tag{11}$$

where z_{iq} is a binary indicator of node is membership in latent class q. This likelihood is the same as the likelihood for a binary logistic regression, and as such can be maximized in θ using standard ML estimation, specifically iteratively weighted least squares.

5.3 A Note on Standard Errors

In Wedel and DeSarbo's [20] description of a mixture model on the general linear model, they note that the standard errors of the estimates can be calculated using a weighted information matrix, taking into account the uncertainty of latent class assignments. However, this relies on the mixture EM approach. In a classification EM-like approach, the parameters are estimated at each maximization step with each node assigned to a class with a weight of 1. As such, the standard errors of the estimates can be interpreted as the standard error of the estimated parameters if the estimated latent class solution was observed and true.

5.4 Estimation for SRFM-ERGMs

There are two additional considerations for estimating SRFM-ERGMs. The first is that of multiple start values. The EM algorithm, either classification or mixture type is susceptible to local maxima [41, 25]. To account for this, we initialize the estimation with multiple start values. Once every model has converged, the latent class solution with the greatest likelihood is selected.

Additionally, there are the multiple issues with using the MPLE to obtain parameter estimates. To account for this, once the EM algorithm has converged the latent class labels are used to estimate an MLE ERGM. The parameter estimates of that model are presented.

The full estimation algorithm proceeds as follows:

1. Initialize starting values for latent class labels $\mathbf{Z}_0$
2. For iteration r
   (a) Maximize the likelihood $\ell(\theta_r|Z_{r-1})$ to obtain θ_r
   (b) Compare $\ell(\theta_r|Z_{r-1})$ against $\ell(\theta_r - 1|Z_{r-2})$. If change in log-likelihood is less than tolerance, estimation has converged to a solution.
   (c) Obtain the hard latent class labels using $E(Z_r|\theta_r)$
3. Once estimation has converged, save solution, initialize a new set of starting values and return to step 2.

4. Once solutions are saved for a pre-defined number of starting values, choose the converged solution with the greatest final $\ell(\theta|Z)$. This solution has the most probable set of latent class labels out of all of the solutions.

5. With the most probable solution selected, fit a MLE ERGM considering the latent class labels as observed, and return the parameter estimates θ.

6 Methods

In order to test the performance of the SRFM-ERGM and to assess the consequences of neglecting to model latent class heterogeneities, three sets of simulations were performed. The first set of simulations demonstrates the consequences of fitting homogeneous ERGM models to networks that are truly heterogeneous. The second set of simulations demonstrates the improvement in recovery of model parameters when using SRFM-ERGMs on heterogeneous networks. Finally, the third set of simulations demonstrates improvement in latent class labels as sample size increases

To more accurately reflect the type of data a researcher would encounter as well as to follow good practice in ERGM methodology development [27], the simulations were based off of an empirical dataset, consisting of a single network of 151 middle schoolers. In addition to directed friendship nomination, information on gender, ethnicity, tobacco use, alcohol use, marijuana use and antisocial behavior were collected. Heterogeneity on sender alcohol use, sender absolute difference in alcohol use, and the edge parameter was modeled in the empirical analysis. For detailed descriptive statistics, and results of the empirical analysis see the Empirical Example section.

Due to missing data at the covariate levels, multiple imputation was used to generate a set of covariate datasets for the empirical analysis in Section 3. In these simulations, a single dataset of covariates was randomly selected from the multiply imputed datasets, and was used for all simulated networks. This follows the simulation procedure outlined in Van Duijn, Giles and Handcock (2009). For each simulation trial, a new network was simulated without changing the covariate dataset.

There were 9 simulation conditions in total (see Table 1). Simulation conditions were arrived at by examining the parameter estimates from the empirical model and changing some of the parameter values to reflect different situations. For example, the empirical model indicated heterogeneity in the effect of both sender alcohol use and the edge term. For condition 1 of the simulations, which only modeled a heterogeneous effect of the edge term, the parameter for sender alcohol use was set to be equal across both classes. Additionally, for all conditions the edge parameter was increased from the empirical example to simulate less sparse matrices and therefore increase the amount of information in the data. Finally, in all conditions, approximately 25% of nodes were in class 2. This was derived from the empirical results. Four out of the 9 simulation conditions were replications of conditions, but with increased effect sizes. These increased effect size conditions are denoted with a +. Increased effect sizes were assessed to get a sense of what approximate effect sizes return accurate latent class labeling.

In the third set of simulations, the size of the network N was manipulated. In an attempt to keep the nodal covariate distributions similar as sample size increased, for each value of N, the node level data was formed by sampling with replacement from the 151 original individuals' data. Once one dataset per value of N was generated, 500 networks using Condition 3 parameter values and

500 networks using Condition 4 parameter values were generated and correctly specified models were fit.

Table 1: Simulation Conditions: Values are the data generating parameters per condition.

Condition	1	1+	2	2+	3	3+	4	4+	5
GWESP (.1) Class 1	0.97	0.97	0.97	0.97	0.97	0.97	1	1.2	0.97
GWESP (.1) Class 2	0.97	0.97	0.97	0.97	0.97	0.97	0.2	0.2	0.97
Mutual	2.34	2.34	2.34	2.34	2.34	2.34	2.34	2.34	2.34
Gender Match	0.52	0.52	0.52	0.52	0.52	0.52	0.52	0.52	0.52
Ethnicity Match	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84
Sender Tobacco Use	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13
Receiver Tobacco Use	-0.11	-0.11	-0.11	-0.11	-0.11	-0.11	-0.11	-0.11	-0.11
Sender Anti Social	-0.64	-0.64	-0.64	-0.64	-0.64	-0.64	-0.64	-0.64	-0.64
Receiver Anti Social	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24
Sender MJ	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33
Receiver MJ	0.21	0.21	0.21	0.21	0.21	0.21	0.21	0.21	0.21
Receiver Alcohol Use	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
Abs Diff Tobacco Use	-0.21	-0.21	-0.21	-0.21	-0.21	-0.21	-0.21	-0.21	-0.21
Abs Diff MJ	-0.39	-0.39	-0.39	-0.39	-0.39	-0.39	-0.39	-0.39	-0.39
Abs Diff Anti Social	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08
Abs Diff Alcohol Use	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
Sender Alcohol Use Class 1	0.06	0.06	0	0	0	0	0.06	0.06	0.06
Sender Alcohol Use Class 2	0.06	0.06	-0.25	-1	-0.25	-1	0.06	0.06	0.06
Edges Class 1	-3.5	-3.25	-3.85	-3.85	-3.5	-3.25	-3.65	-3.65	-3.85
Edges Class 2	-4.25	-4.5	-3.85	-3.85	-4.25	-4.75	-3.65	-3.65	-3.85

MJ: Marijuana, Abs Diff: Absolute Difference

In Conditions 1 through 4 and 1a though 4a, correctly specified models were fit to the data as well as homogeneous models. In Condition 5 a model that specified 2 latent classes and heterogeneity on the sender alcohol use parameter and the edges parameter was estimated, along with a homogeneous model.

Raw and relative biases were assessed, as well as computed standard errors from the simulation set and average estimated standard errors. Additionally, Hubert-Arabie Adjusted Rand Indices (ARI; [42, 43]) were computed for the comparison of the estimated latent class labels and the true latent class labels. This indicates performance of the method, in that if the method failed to detect true latent class labels for a condition, the method would have not performed well with that parameter set. The Adjusted Rand Index is a measure of agreement that accounts for the expected level of agreement if the labels were assigned randomly. Finally, for Conditions 2 and 2+ Adjusted Rand are presented for the latent class labels of only the nodes that had a value higher than 0 for alcohol use. By examining the subset of nodes with positive alcohol use, we can examine the performance of the method when the latent classes are defined only for a subset of nodes.

All conditions had 500 networks simulated. The networks were simulated using the MCMC

approach used in the R package statnet [44]. This approach randomly toggles edges within a network with the probability according to the generating model. This allows the simulated data to properly reflect transitivity, as well as contain sampling variability. To clarify, only the networks themselves were simulated. The covariate data set was the same across all simulations as was the true latent class labeling. Estimation of these models used twenty random start values per trial as per the algorithm described above.

7 Results

7.1 Parameter Recovery for Homogeneous Models

What follows are the results from fitting homogeneous models to the data generated from the conditions. Relative biases are presented in Table 2. Tables of mean parameter value estimates, raw bias and estimates of standard errors are contained in Appendix A. Note that for Conditions 1 through 4+, the homogeneous model fitted is not the generative model, while for Condition 5, the homogeneous model is the true model.

Table 2: Relative Bias for Homogeneous Models calculated with Mean Raw Bias over True Parameter Values. Results indicate that there is a substantial and systematic pattern of bias in the homogeneous parameters when heterogeneity is not properly modeled.

Condition	1	1+	2	2+	3	3+	4	4+	5
GWESP (.1) Class 1	**0.131**	**0.330**	0.024	**0.102**	**0.177**	**0.672**	*-0.465*	*-0.512*	0.037
GWESP (.1) Class 2	-	-	-	-	-	-	*1.675*	*1.930*	-
Mutual	0.003	0.008	0.000	-0.004	0.003	-0.003	0.002	0.001	0.002
Gender Match	-0.029	-0.077	-0.006	-0.017	-0.040	**-0.154**	-0.073	-0.088	-0.002
Ethnicity Match	-0.065	**-0.155**	0.005	-0.040	-0.087	**-0.263**	**-0.148**	**-0.176**	0.014
Sender Tobacco Use	**0.392**	**0.608**	**0.208**	**0.531**	**0.608**	**1.085**	**0.385**	**0.485**	**-0.100**
Receiver Tobacco Use	**-0.118**	**-0.209**	0.009	-0.091	**-0.182**	**-0.436**	**-0.145**	**-0.200**	**0.164**
Sender Anti Social	**-0.458**	**-0.783**	**-0.158**	**-0.363**	**-0.572**	**-1.002**	**-0.688**	**-0.823**	-0.013
Receiver Anti Social	**-0.108**	**-0.175**	-0.042	**-0.129**	-0.088	**-0.367**	**-0.138**	**-0.188**	-0.017
Sender MJ	**-0.291**	**-0.524**	-0.094	-0.085	**-0.415**	**-0.567**	**-0.430**	**-0.530**	-0.009
Receiver MJ	-0.033	-0.043	0.019	-0.010	-0.057	-0.043	-0.005	-0.014	-0.005
Receiver Alcohol Use	**-0.100**	**-0.450**	**-0.250**	**-1.275**	**-0.350**	**-0.650**	**-0.375**	**-0.525**	-0.025
Abs Diff Tobacco Use	0.038	0.019	-0.010	0.010	0.043	0.024	-0.048	-0.038	-0.076
Abs Diff MJ	-0.018	-0.064	0.013	-0.003	-0.031	-0.108	-0.046	-0.064	-0.005
Abs Diff Anti Social	**-1.525**	**-2.388**	**-0.638**	**-1.913**	**-2.113**	**-3.688**	**-2.063**	**-2.413**	-0.025
Abs Diff Alcohol Use	**0.100**	**-0.200**	**-0.800**	**-5.100**	**-1.300**	**-3.400**	**0.100**	**-0.100**	**0.100**
Sender Alcohol Use Class 1	**0.817**	**1.333**	-	-	-	-	**0.983**	**1.250**	-0.017
Sender Alcohol Use Class 2	-	-	*-0.416*	*-0.533*	*-0.728*	*-0.836*	-	-	-
Edges Class 1	*0.219*	*0.421*	0.016	0.055	*0.244*	*0.566*	0,055	0.053	0.012
Edges Class 2	0.004	0.026	-	-	0.024	0.071	-	-	-

MJ: Marijuana, Abs Diff: Absolute Difference
Bold text indicates relative bias in homogeneous parameters above .1, *italic text* indicate relative bias in heterogeneous parameters above .1.

As expected, all of the heterogeneous conditions (1 through 4+) demonstrated large relative bias

in many of the homogeneous parameters (Highlighted in Table 2 in bold) and increased raw bias (See Table A4). Five effects that had consistent relative bias above .1 in all heterogeneous conditions were Sender Tobacco Use, Sender Anti Social, Receiver Alcohol Use, Absolute Difference in Anti Social and Absolute Difference in Alcohol Use. It should be noted that the absolute Difference in Anti Social and Absolute Difference in Alcohol Use both had true parameters close to 0 (-.08 and -.1 respectively), therefore a certain amount of relative bias is expected. As shown in Table A4, there is increased raw bias in these parameter estimates as one increases the difference between the latent classes.

Another result of note is the pattern of bias in the GWESP term for Simulations 1 through 3+. With the exception of Condition 2, the relative biases in the estimate of the GWESP term are above .1. This suggests that the GWESP term is sensitive to misspecification of other heterogeneous effects as homogeneous, and it responds to misspecification by increasing in magnitude. Intuitively, this is due to transitivity effects being intrinsically confounded with covariate effects. An edge might form because of a similarity on behavior, or due to the influence of shared partners, and if the effect of similarity on behavior is underestimated, the effect of shared partners might be overestimated to compensate. This bias in a structural effect is not replicated in the mutuality parameter, which suggests that mutuality is robust to misspecifications in the rest of the model.

As expected, there was bias in the parameter estimates that were truly heterogeneous (italicized in Table 2), as it would be impossible to correctly recover the true parameter estimates that were heterogeneous with a single parameter that is homogeneous. Note no relative bias was calculated for the Sender Alcohol Use of Class 1 for Conditions 2 through 3+ as the true value of the parameter was 0. As for Condition 5, while there was some cases of relative bias being higher than .1 (Sender and Receiver Tobacco Use, and Absolute Difference in Alcohol Use) the rest of the parameters were recovered well. This was expected, as for Condition 5 a homogeneous model is the true model.

In summary, Tables 2 suggest that fitting homogeneous models to networks that are heterogeneous results in unpredictable patterns of high bias in both parameters that are supposed to be heterogeneous, as well as parameters that are indeed homogeneous.

7.2 Mixture Component Label Recovery with SRFM-ERGMs

Table 3 shows the mean Adjusted Rand Indices for latent class recovery for every individual, as well as individuals who use alcohol. Table 4 contains the relative bias. As described below, these simulations indicate that the model can recover true latent class structure, and gives an indication as to what level of class difference would be necessary. Tables of mean parameter value estimates, raw bias and estimates of standard errors are contained in Appendix A.

Table 3: Mean Adjusted Rand Indices (Standard Deviations computed from the sample) computed across 500 trials within each condition. Results indicate best recovery occurs in the + conditions, with higher class separation. Structural effects such as GWESP are also recovered well.

Condition	ARI	ARI Alcohol Users
1	0.541 (.126)	0.532 (.179)
1+	0.924 (.051)	0.89 (.104)
2	0.06 (.120)	0.272 (.137)
2+	0.1 (.178)	0.703 (.108)
3	0.629 (.120)	0.798 (.137)
3+	0.959 (.034)	0.997 (.014)
4	0.667 (.096)	0.678 (.153)
4+	0.862 (.069)	0.848 (.113)

Adjusted Rand Indices (ARI) were computed as the empirical mean and empirical standard deviation over every trial in a given condition. For class recovery, the mean ARI contained in Table 3 indicate that across all simulation conditions, increasing the difference in parameter weights between the latent classes leads to increased recovery of the latent classes. However, for conditions with smaller class separation (Conditions 1, 2, 3, 4), the recovery of the latent classes was not optimal. For most conditions, the greatest improvement in recovery of the latent class labels occurred only when looking at alcohol users. The exception being conditions 2 and 2+ where even larger differences between the latent classes, recovery of the class labels for alcohol users is still not as good as for other conditions (ARI of .703 for Condition 2+). This is likely due to the latent class being defined solely on individuals who used alcohol, a subset of the sample, as opposed to other conditions that had the latent classes also be defined by structural parameters such as edges or GWESP. These results indicates caution when using SRFM-ERGM, as the effect size for a latent class that is wholly defined by a covariate needs to be larger than the effect size for a latent class that is defined by covariates and network structures.

One set of results of note is that Condition 4 and 4+ had remarkably good recovery of the latent class labels, even at the lower class difference (ARI of .667 and .862 for Condition 4 and 4+ respectively). This suggests that heterogeneities in the GWESP term are reflected strongly in observed network structure. Finally, standard deviations of the ARIs suggest that there was variability in recovery rates due to sampling fluctuations. This variability can lead, particularly in conditions with low class separation, to very poor rates of classification. As one would expect, there is more variability in the ARIs the lower the average ARI becomes.

7.3 Recovery of Parameters in Heterogeneous Data using SRFM-ERGMs

Table 4: Relative Bias in the Parameter Estimates For Heterogeneous Models, calculated using Mean Raw Bias over True Parameter Value. Results indicate modest levels of bias in homogeneous parameters. This bias is in line with expected bias from the approximate MLE estimation

Condition	1	1+	2	2+	3	3+	4	4+	5
GWESP (.1) Class 1	0.027	0.009	0.013	0.025	0.030	0.010	*-0.108*	-0.091	0.048
GWESP (.1) Class 2	-	-	-	-	-	-	*0.107*	*0.228*	-
Mutual	0.000	0.000	0.000	-0.001	-0.001	-0.001	0.000	0.000	0.001
Gender Match	-0.002	0.006	-0.003	0.002	-0.003	0.001	-0.013	-0.016	0.003
Ethnicity Match	-0.004	0.010	0.005	-0.002	-0.005	0.014	-0.023	-0.025	0.017
Sender Tobacco Use	**0.332**	**0.340**	-0.013	-0.088	**0.175**	0.032	**0.191**	**0.240**	-0.035
Receiver Tobacco Use	0.032	0.099	**0.131**	**0.198**	**0.145**	**0.188**	0.063	0.060	**0.167**
Sender Anti Social	**-0.285**	-0.071	**0.108**	**0.227**	**-0.201**	-0.031	**-0.272**	**-0.182**	-0.072
Receiver Anti Social	-0.059	-0.052	0.007	0.049	0.007	-0.051	-0.034	-0.058	-0.012
Sender MJ	**-0.180**	-0.067	-0.090	0.017	**-0.196**	-0.071	**-0.230**	**-0.216**	0.016
Receiver MJ	-0.010	0.010	0.029	0.010	-0.043	-0.065	0.007	-0.003	0.009
Receiver Alcohol Use	0.083	-0.011	**-0.145**	**-0.263**	-0.012	**0.198**	-0.036	**-0.100**	-0.013
Abs Diff Tobacco Use	-0.013	-0.039	-0.066	-0.100	-0.073	-0.093	-0.051	-0.054	-0.074
Abs Diff MJ	0.004	0.002	0.019	-0.001	-0.014	-0.026	0.002	-0.009	0.003
Abs Diff Anti Social	**-0.425**	**-0.165**	0.085	**0.517**	**-0.353**	0.011	**-0.495**	**-0.486**	-0.032
Abs Diff Alcohol Use	**0.341**	**0.137**	**-0.570**	**-0.950**	-0.055	**0.638**	**0.232**	0.036	**0.141**
Sender Alcohol Use Class 1	**0.291**	0.075	NA	NA	NA	NA	**0.502**	**0.388**	**-0.110**
Sender Alcohol Use Class 2	-	-	0.050	-0.013	*-0.163*	-0.003	-	-	**-0.237**
Edges Class 1	0.055	0.015	-0.001	-0.004	0.049	0.008	0.020	0.009	-0.020
Edges Class 2	0.031	0.006	-	-	0.028	-0.049	-	-	0.064

MJ: Marijuana, Abs Diff: Absolute Difference, NAs due to true effect being 0.
Bold text indicates relative bias in homogeneous parameters above .1, ***italic text*** indicates relative bias in heterogeneous parameters above .1.

In terms of bias in the parameter estimates, recovery of heterogeneous parameters was far superior than seen in the results for the homogeneous models discussed above. Importantly, the homogeneous effects (such as gender match, ethnicity match, etc.) also evidenced less bias (see Table A-2). Indeed, very few parameters had greater than .1 absolute raw bias or greater than .1 relative bias.

While modeling heterogeneity greatly improved the degree of bias, there were a few parameters that still evidenced suboptimal parameter estimation. The greatest bias was seen in Conditions 4 and 4+, where the relative bias for the recovery of the lower GWESP term (GWESP Class 2) had a true value of .2 in both conditions. The relative bias here was .107 for Condition 4 and .228 for Condition 4+. However, the raw bias (contained in Supplementary Material Table A2) in terms of magnitude was quite small, .021 and .046 respectively. This level of bias is not relevant to researchers analyzing empirical data, and is likely due to the GWESP term being difficult to estimate in general. Of note, the GWESP estimates contained very little relative or absolute bias in the other conditions once heterogeneity was modeled.

The one other heterogeneous term that had a relative bias of above .1 was that of Sender Alcohol Use in Class 2 for Condition 3. The relative bias for this parameter was -.163, with a raw bias of

.041. The true parameter value was -.25, which suggests that while the relative bias was high, the actual level of raw bias was comparatively low. Additionally, compared to the homogeneous results (where 4 conditions evidenced high bias on this effect), that only one condition in the heterogeneous modeling framework evidenced bias is a large gain.

A few homogeneous effects contained bias. The Sender Anti Social effect for conditions 1, 2+, 3, 4, and 4+ all have relative bias magnitude greater than .1. This effect was held homogeneous across all conditions and had the highest magnitude of the homogeneous effects (-.64). The patterning of bias across all conditions was not consistent, and this likely indicates that the effect was being disrupted by the effects of the structural parameters such as GWESP.

Additionally, with Sender Tobacco Use, several conditions had relative bias greater than .1. However, the true effect of Sender Tobacco Use was small (.13) in all conditions, which can lead to small levels of raw bias translating into large relative bias. This occurrence of high relative biases for very small true parameters is expected and most notably occurs for the Absolute Difference of Anti-Social Behavior, Sender Alcohol use (for models with homogeneous Sender Alcohol Use), and Absolute Difference in Alcohol Use. Furthermore, for conditions 2, 2+, 3 and 3+ the true effect of Sender Alcohol Use in Class 1 was 0, which would lead to undefined relative bias for any level of raw bias.

As for Condition 5, recovery of homogeneous parameters was quite good, with no raw bias being greater than .1, and very few relative biases greater than .1. This suggests that the misspecification of the latent class structure did not lead to systematic bias in the homogeneous parameters. Additionally, the spread of the latent class parameters around the true homogeneous parameters was quite reasonable, with Sender Alcohol Use having a true effect of .06, and the latent classes returning an effect of .053 and .046, while the Edge effect is 3.85 while the latent classes returned effects of 3.773 and 4.097 respectively.

In summary, Table 4 suggest that the sender latent class models fit to heterogeneous data do a reasonable job of recovering the model parameters. Importantly, the results are far less biased than those seen in the results when homogeneous models are fit to the heterogeneous data sets. Bias in the parameter estimates is in part due to ERG modeling being approximate even when using the MCMC-MLE. A degree of bias is to be expected. That being said, the fact that bias still remained even when fitting the generating model to data suggests that researchers should have stringent criteria for interpreting results from latent class ERGMs, interpreting both effect size as well as significance level.

As for estimated standard errors of the estimates and the standard deviations across the simulation trials for every homogeneous parameter, the estimated standard error and the standard deviation of the parameter across the condition agreed. However, for the heterogeneous parameters the estimated standard error is considerably smaller than the simulation standard deviation (Table A6 in supplementary material). This is in part due to error in classification across the simulation trials. If there are simulation trials where the individuals are misclassified then the estimates of the heterogeneous effects will be more variable than if in every trial the latent class labels were recovered with perfect fidelity. Additionally, the estimated standard errors assume that the latent class assignment is the true assignment, and are not adjusted for uncertainty in class assignment. This is a reflection of the classification likelihood EM algorithm in use in the estimation.

7.4 Sample Size Simulations

Figure 1 and 2 shows the distribution of Adjusted Rand Indices as a function of increased sample size. Classification accuracy improves as one would expect as a function of sample size. The lack of a smooth curve in classification accuracy is a result of the simulation procedures. While 500 networks were simulated per sample size target, only one covariate set was generated for a given sample size. This allows the variation in classification to be a function mainly of the network, rather than the covariate set, however results in some counter intuitive results in the classification, such as the jumps in classification accuracy exhibited for several sample sizes. The overall shape of the curve suggests that the estimates of the latent class labels are improving with sample size.

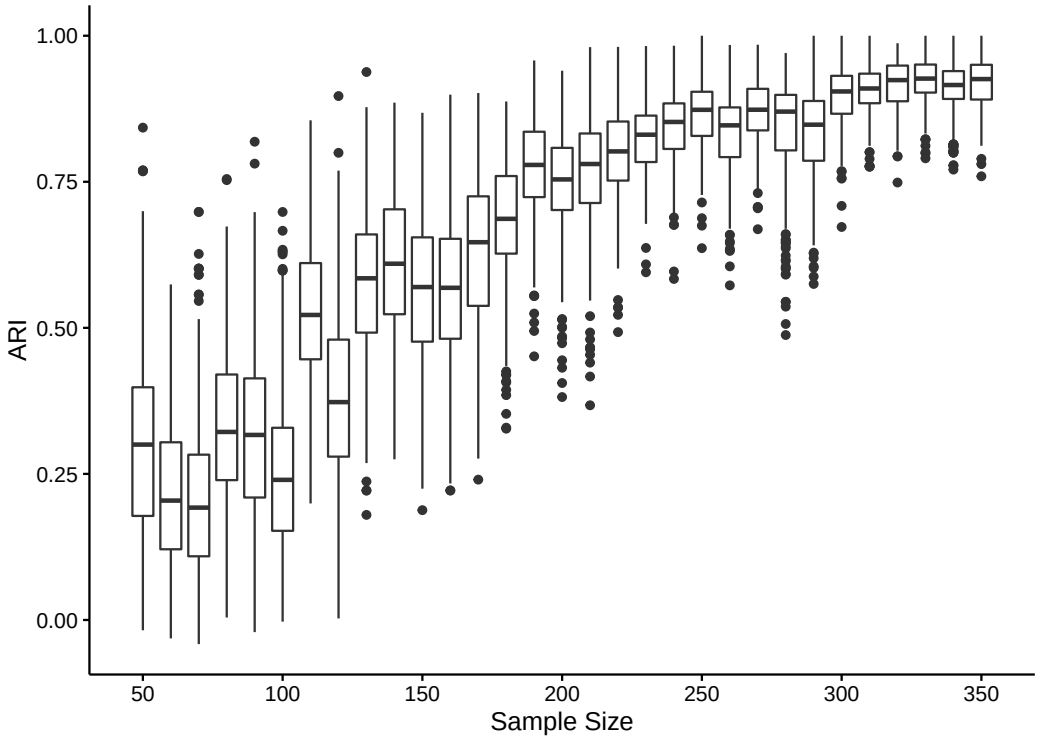

Figure 1: ARI by Sample Size for Condition 3 simulations

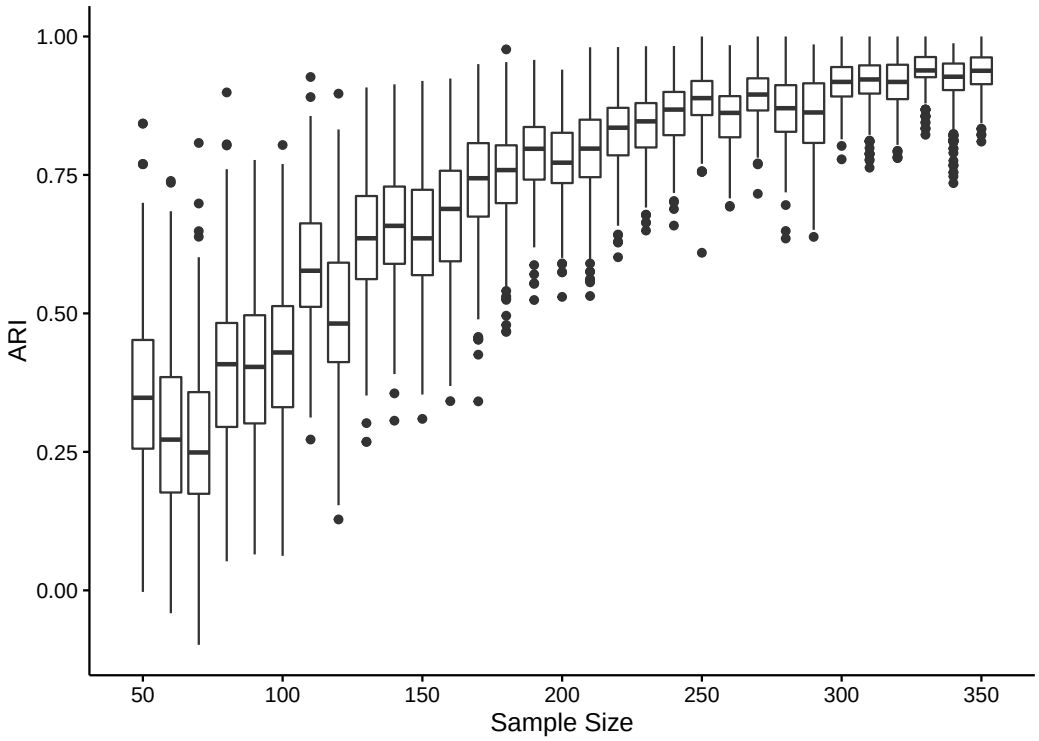

Figure 2: ARI by Sample Size for Condition 4 simulations

8 Empirical Example

The above simulations were based on an empirical data set of high middle school students. In this empirical example, we examine potential heterogeneity in the effect of alcohol use on peer assortment. Peer assortment on the basis of substance use and other deviant behaviors has been a continually studied topic in clinical psychology [10], however, whenever network models are applied to network data to study peer assortment on substance use, researchers have used homogeneous models. Heterogeneous effects of alcohol use on peer assortment could potentially have important implications for network intervention design, as an intervention model based on homogeneous results may not be applicable to all at risk individuals in the intervention population.

We use alcohol use here due to its fairly high use rate among this sample (26.67% of the sample have used alcohol). In future studies of heterogeneous peer assortment and selection effects of substance use, other substances such as marijuana and tobacco should be considered.

Three parameters were used to model each of the effects of alcohol use, marijuana use, tobacco use and anti-social behavior. Indegree parameters measure the effect of the variable on an individuals tendency to receive friendship nominations, outdegree parameters measure the effect of a variable on an individuals tendency to send friendship nominations. Finally, absolute difference terms measure the effect of different values of the variable on the tendency for two individuals to

form a friendship.

In the heterogeneous model, we propose a sender-type heterogeneity for the outdegree effect of alcohol use and the absolute difference effect of alcohol use. This can be interpreted as allowing the effect of alcohol use to differ in its effect on an individuals tendency to nominate friends both generally, and with regards to the potential friends alcohol use.

9 Data Characteristics and Measures

This empirical data-set consisted of a 151 students who attended the same high school. Students were assessed on a variety of demographics and risk behaviors. The current analysis focuses on substance use of three types, alcohol, tobacco and marijuana, as well as general anti-social behavior. Descriptive statistics of the data are provided in Table 5.

Table 5: Demographic and Descriptive Statistics

	Male	Female	Missing			
Gender	46%	51.30%	2.60%			
	African-American	Asian	Latino/Hispanic	White	Mixed	Missing
Ethnicity	23.03%	1.32%	18.42%	48.03%	6.58%	2.63%
	Mean	25% Quantile	Median	75% Quantile	% Above 1	Missing
Age	15.07	14.65	15.01	15.33	-	0.66%
Alcohol Use	0.4476	0	0	1	26.97%	5.92%
Tobacco Use	0.1724	0	0	0	6.58%	4.61%
Marijuana Use	0.1781	0	0	0	9.21%	3.95%
Anti social behavior	0.305	0.129	0.258	0.4516	-	4.61%

9.1 Measures

Friendship Nominations. Individuals were provided a full roster of the school and were allowed to nominate any number of individuals as friends. These nominations were binary (Yes a friend, or No not a friend), and directed.

Alcohol Use. Alcohol use was assessed with a single item which asked "In the last 6 months on how many days did you have at least one drink of alcohol?" With the following answer choices: 1) 0 days, 2) 1-2 days, 3) 3-5 days, 4) 6-9 days and 5) 10 or more days.

Tobacco Use. Tobacco use was assessed with a single item which asked "In the last 6 months on how many cigarettes did you smoke per day?" With the following answer choices: 1) 0, 2) 1, 3) 2-3 , 4) 6-10 and 5) 10-20, 6) more than a pack.

Marijuana Use. Marijuana use was assessed with a single item which asked "In the last 6 months how many times did you use marijuana?" With the following answer choices: 1) 0, 2) 1-2 times, 3) 3-9 times, 4) 10-19 and 5) 20 or more times.

Anti Social Behavior. Antisocial behavior was assessed with a subset of 31 items from the Youth Self Report scale [45]. These items assessed behaviors over the previous 6 month period and had text such as "I get in many fights", "I set fires," and "I tease others a lot." Responses were on a 3 point scale with 0 being "Not True", 1 being "Somewhat True", and 2 being "Very/Often True." These items were mean scored for a final Anti-Social behavior Composite.

9.2 Missing Data

Due to the inability for ERGMs to handle missing data at the nodal attribute level, multiple imputation was used to account for missingness [46]. 500 covariate datasets were simulated using information on gender, age, ethnicity, substance use and all anti social behavior items. Results for the 500 covariate datasets were combined according to Rubin (1987). There are two competing interests to the use of multiple imputation here. As was mentioned, ERGMs cannot account for missingness among the nodal covariates, as such one option would be to remove those individuals from the network. However, as networks are highly interdependent systems by definition, this would substantially change the network structure, possibly leading to different results. Multiple imputation however, does not take into a account the dependency between individuals on their nodal attributes, rather assuming them to be independent. This assumption leads to an attenuation of any homophily effects within the network. In the case of this application, it was decided that the damage to the network structure outweighed the attenuation of homophily, and multiple imputation was used.

9.3 Model Specification

Of substantive interest for these analysis is potential heterogeneity in the sender effects of alcohol use, as well as the Sender Effect of the Absolute Difference in Alcohol Use. Additionally, a general heterogeneity was modeled using a heterogeneous edge parameter. This model nearly identical to the model presented in Condition 3 of the simulations was fit. The sole addition to the model was allowing the effect of Absolute Difference in Alcohol Use to vary between the two latent classes. This heterogeneous effect was not included in the simulation conditions due to a lack of heterogeneity in the effect exhibited in the empirical data, as shown below. Additionally, a homogeneous model was fit to the empirical data.

10 Results of the Empirical Analysis.

Table 6: Empirical Results for both heterogeneous model and homogeneous model. Estimates are in mean log-odds metric calculated across 500 multiply imputed datasets with (multiple imputation standard error)

Parameter	Heterogeneous Model	Homogeneous Model
GWESP (alpha = .1)	0.97 (0.07)*	1.25 (0.07)*
Mutuality	2.34 (0.14)*	2.26 (0.15)*
Gender Match	0.52 (0.06)*	0.44 (0.06)*
Ethnicity Match	0.85 (0.06)*	0.47 (0.06)*
Sender Tobacco Use	0.13 (0.11)	0.29 (0.10)*
Receiver Tobacco Use	-0.11 (0.10)	-0.05 (0.09)
Sender Anti Social	-0.64 (0.21)*	-0.23 (0.20)
Receiver Anti Social	0.24 (0.20)	0.24 (0.20)
Sender MJ	0.33 (0.11)*	0.14 (0.10)
Receiver MJ Use	0.21 (0.11)	0.19 (0.10)
Receiver Alcohol Use	0.04 (0.07)	-0.02 (0.06)
Abs Difference Anti Social	-0.08 (0.11)	-0.02 (0.18)
Abs Difference Tobacco Use	-0.21 (0.10)*	-0.13 (0.10)
Abs Difference MJ Use	-0.39 (0.19)*	-0.18 (0.10)
Abs Difference Alcohol Use Class 1	-0.01 (0.08)	-0.01 (0.06)
Abs Difference Alcohol Use Class 2	-0.05 (0.07)	
Sender Alcohol Use Class 1	-0.21 (0.10)*	0.02 (0.06)
Sender Alcohol Use Class 2	0.06 (0.07)	
Edges Class 1	-5.28 (0.11)*	-5.09 (0.10)*
Edges Class 2	-4.22 (0.12)*	

MJ: Marijuana, Abs Difference: Absolute Difference
*: p <.05

Table 6 contains the results for the heterogeneous and homogeneous analysis. Of note in these results is the effect of Sender Alcohol Use in Class 1 versus the effect of Alcohol Use in Class 2 . These effects are significantly different from one another (z = -2.21, p <.05), suggesting that for members of class 1, alcohol use reduces the number of friendship nominations sent out while for members of Class 2, alcohol use does not effect the number of friendship nominations sent out. Additionally, there is a significant difference between the edges effect of Class 1 and Class 2 (z = 6.51, p <.05). This suggests that members of Class 1 have fewer friendship nominations made in general, while members of Class 2 make more friendship nominations.

There are several significant homogeneous effects in the heterogeneous model. The effect of Sender Anti Social Behavior is negative suggesting that individuals with more antisocial behavior nominate fewer friends. The effects of Absolute Difference in Tobacco and Marijuana Use are

both negative suggesting that individuals who are different on their usage of those substances are less likely to nominate each other as friends. Finally there was a positive significant effect of Sender Marijuana Use suggesting that individuals who use more Marijuana make more friendship nominations in general.

There are expected positive significant findings for GWESP, mutuality, gender and ethnicity matching. The significant GWESP finding suggests that individuals who share friends are more likely to be friends themselves. Additionally there is a strong effect for mutuality suggesting that reciprocal friendship nominations are the norm in this network.

As for the differences between the heterogeneous model and the homogeneous model, it is apparent that the modeling of heterogeneous effects of alcohol use and the edge parameter clarifies several of the homogeneous parameters. Specifically, the effects for Sender Anti Social Behavior, Sender Marijuana Use, Absolute Difference in Tobacco Use and Absolute Difference in Marijuana Use are significant in the Heterogeneous Model and not significant in the Homogeneous Model. Interestingly, one effect, that of Sender Tobacco Use, was significant in the homogeneous model and not significant in the heterogeneous model.

It is worthwhile to note that significantly fewer individuals were classified into class 2 than into class 1, with the average proportion of individuals in class 2 across all the multiply imputed datasets being 27.8%. Class separation on the basis of class membership probability was exceptionally good for every multiply imputed dataset. The average probability of membership in class 1 for individuals with a higher than .5 probability of being in class 1 was .939, while the average probability of class membership in class 2 for individuals with a higher than .5 probability of being in class 2 was .9671. This level of class separation is in part due to the classification EM algorithm used in estimation, which tends to find the latent class labeling that is the most informative rather than the latent class probabilities that maximize the likelihood.

In summary, the findings support a level of heterogeneity in the effect of alcohol use, specifically for the effect of alcohol use on general outdegree. Indeed, some individuals are less likely to nominate friends the more alcohol they drink, while other individuals exhibit no effect of alcohol use on friendship formation. Of broader interest, there are more significant findings in the heterogeneous model as compared to the homogeneous model. This suggests that the inclusion of the heterogeneous model parameters clarified the effects of the other parameters.

11 Discussion

This manuscript presented the Sender/Receiver Latent Class ERGM, provided a estimation routine, presented simulations that demonstrated its ability to detect latent class structures in networks, and demonstrated the consequences of failing to model latent class structures in network. Finally, this manuscript presented an empirical example rooted in adolescent peer networks that demonstrated this models ability to obtain reasonable results from real, rather than simulated data.

There are several broad results that can be gleaned from the simulations. First, recovery of the latent class structures is sensitive to the difference between the latent classes, as demonstrated by the increased difference conditions in the simulations. This is to be expected; however, caution is advised when using these models on empirical data as the models will return solutions regardless

of underlying differences between the latent classes. One solution to this problem is examining the *a posteriori* probabilities of class membership given the returned set of parameter estimates. If there is a large amount of uncertainty in the classifications of individuals, then the model should be reassessed. Simulation results suggest that the classification probabilities should be quite high for correctly classified individuals. Additionally, penalized likelihood criteria, such as AIC or BIC, or entropy based criteria (e.g. NEC; [47, 48], Conditional Classification Criteria; [49]), can provide guidance for determining the presence and the number of latent classes in empirical data.

The main point that can be taken away from this study is that caution is warranted when analyzing data with homogeneous models. As shown in the simulations, mis-specifying structural parameters such as edges or GWESP can result in wide-ranging bias of many other model parameters, where this bias is generally unpredictable in its direction. Additional caution is warranted when looking at the effects of covariates as a mis-specification of one covariate effect (such as a sender effect) as homogeneous when really there are heterogeneous effects leads to bias in other effects of the same covariate (such as absolute difference). This is particularly concerning as the direction of the bias is unpredictable, and therefore mis-specification could lead to spurious findings.

As another finding of particular interest, mis-specifying a heterogeneous model to a truly homogeneous network did not result in any undue bias, other than in the mis-specified terms. Even with the mis-specified terms the differences in effect size were small. This is in part due to the nature of the classification likelihood, in that it tends to collapse components that are ill separated. This suggests that researchers can use heterogeneous models as an exploratory method, without worry that a heterogeneous mis-specification of a homogeneous network would result in biased estimates.

The empirical example offers support for this method in answering substantive questions. Here, that question was: does the effect of alcohol use on the number of friends one nominates differ between individuals? With this sample, there appears to be at least some level of heterogeneity in the effect of alcohol use on the number of friends an individual nominates. Furthermore, by modeling heterogeneity, several effects were clarified that would have not been significant if a homogeneous model was run. These findings further emphasize the importance of modeling individual differences in network data.

There were several limitations to the Sender/Receiver latent class ERGMs. First, the Sender/Receiver Latent Class ERGM requires a directed network. In many cases, the network, and the relation that defines the edges in the network would not be directed, and these models would not be able to be fit. Secondly, these models require the analyst to make a choice to model either Sender latent classes, or Receiver latent classes, but not both. This restricts the type of heterogeneity that an analyst can model. Finally, as developed the Sender/Receiver Latent Class ERGM is only for cross-sectional data.

However, these limitations open up several new directions for future work. The next step would be extending these models to fit general finite mixture ERGMs, rather than restricting the latent class definition to be sender or receiver. This likely would require several approximations, and therefore a Bayesian approach might be warranted. A general finite mixture ERGM would resemble the ERGM models proposed by Schweinberger and Handcock [16], however, would not focus on latent community detection (though they would be capable of it). Additionally, a general latent class ERGM would be able to model undirected networks. Another extension would be to develop a longitudinal version. Adapting the separable temporal ERGMs [2] which can model

peer selection processes over time, would allow analysts to investigate heterogeneous peer selection processes.

Overall, this manuscript introduces a new way of modeling network data and a different definition of heterogeneity on networks that could better serve substantive researchers studying a variety of heterogeneous and latent network-rooted behaviors and traits, such as in adolescent peer influence of substance abuse (eg., Brechwald & Prinstein, 2011), in clinical domains such as autism research (eg., [50]), and also in social, cognitive, and affective neuroscience [51]. Further work is needed to refine and expand the latent class ERGM so that psychological and social scientists can advance our understanding of ideographic variation in network settings.

References

[1] Tom A.B. Snijders, Gerhard G. van de Bunt, and Christian E.G. Steglich. Introduction to stochastic actor-based models for network dynamics. *Social Networks*, 32(1):44–60, 2010.

[2] Pavel N. Krivitsky and Mark S. Handcock. A separable model for dynamic networks. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 76(1):29–46, 2014.

[3] O Frank and David Strauss. Markov graphs. *Journal of the American Statistical Association*, 81(395):832–842, 1986.

[4] Stanley Wasserman and Philippa Pattison. Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p^*. *Psychometrika*, 61(3):401–425, 1996.

[5] Peter C M Molenaar. A Manifesto on Psychology as Idiographic Science: Bringing the Person Back Into Scientific Psychology, This Time Forever. *Measurement: Interdisciplinary Research & Perspective*, 2(4):201–218, 2004.

[6] D J Bauer. Evaluating Individual Differences in Psychological Processes. *Current Directions in Psychological Science*, 20(2):115–118, 2011.

[7] J Jaccard, R Turrisi, and C K Wan. Interaction effects in multiple regression. *Sage University Paper Series on Quantitative Applications in the Social Sciences*, (07-072), 1990.

[8] D J Bauer and L Cai. Consequences of Unmodeled Nonlinear Effects in Multilevel Models. *Journal of Educational and Behavioral Statistics*, 34(1):97–114, 2008.

[9] S Chatterjee and P Diaconis. Estimating and understanding exponential random graph models. *Annals of Statistics*, 2013.

[10] Whitney A Brechwald and Mitchell J Prinstein. Beyond homophily: A decade of advances in understanding peer influence processes. *Journal of Research on Adolescence*, 21(1):166–179, 2011.

[11] Noona Kiuru, William J. Burk, Brett Laursen, Katariina Salmela-Aro, and Jari Erik Nurmi. Pressure to drink but not to smoke: Disentangling selection and socialization in adolescent peer networks and peer groups. *Journal of Adolescence*, 33(6):801–812, 2010.

[12] Krzysztof Nowicki and Tom A. B. Snijders. Estimation and Prediction for Stochastic Blockstructures, 2001.

[13] Peter D Hoff, Adrian E Raftery, and Mark S Handcock. Latent Space Approaches to Social Network Analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.

[14] Mark S Handcock and Adrian E Raftery. Model-based clustering for social networks. *Journal of the Royal Statistical Society, Series B*, pages 1–22, 2007.

[15] J-J Daudin, Franck Picard, and Stephane Robin. A mixture model for random graphs. *Stat. Comput.*, 18(2):173–183, 2008.

[16] Michael Schweinberger and Mark S Handcock. Local dependence in random graph models: characterization, properties and statistical inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(3):647–676, 2015.

[17] T. M. Sweet. Incorporating Covariates Into Stochastic Blockmodels. *Journal of Educational and Behavioral Statistics*, 2015.

[18] Douglas Steinley, Michael J. Brusco, and Stanley Wasserman. Clusterwise p* models for social network analysis. *Statistical Analysis and Data Mining*, 4(5):487–496, oct 2011.

[19] Wayne S. DeSarbo and William L Cron. A maximum likelihood methodology for clusterwise linear regression. *Journal of Classification*, 5(2):249–282, 1988.

[20] Michel Wedel and Wayne S DeSarbo. A mixture likelihood approach for generalized linear models. *Journal of Classification*, 12(1):21–55, 1995.

[21] PW Holland and S Leinhardt. An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Associationerican Statistical ...*, 76(373):33–50, 1981.

[22] David R Hunter and Mark S Handcock. Inference in Curved Exponential Family Models for Networks. 15(3):565–583, 2006.

[23] Michael Kearns, Yishay Mansour, and Andrew Y Ng. An Information-Theoretic Analysis of Hard and Soft Assignment Methods for Clustering. *Proc. of Conference on Uncertainty in Artificial Intelligence*, pages 282–293, 1997.

[24] J B MacQueen. Some Methods for classification and Analysis of Multivariate Observations. *Proc. Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967.

[25] Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. B*, 39(1):1–38, 1977.

[26] David R Hunter, Steven M Goodreau, and Mark S Handcock. Goodness of Fit of Social Network Models. 103(481):248–258, 2008.

[27] Marijtje A J van Duijn, Krista J Gile, and Mark S Handcock. A Framework for the Comparison of Maximum Pseudo Likelihood and Maximum Likelihood Estimation of Exponential Family Random Graph Models. *Social Networks*, 31(1):52–62, 2009.

[28] Bradley Efron. The Geometry of Exponential Families. *The Annals of Statistics*, 6(2):362–376, 1978.

[29] Garry Robins, Peter Elliott, and Philippa Pattison. Network models for social selection processes. *Social Networks*, 23(1):1–30, 2001.

[30] Miranda J. Lubbers and Tom A B Snijders. A comparison of various approaches to the exponential random graph model: A reanalysis of 102 student networks in school classes. *Social Networks*, 29(4):489–507, 2007.

[31] Francis Comets and Martin Janžura. A central limit theorem for conditionally centred random fields with an application to Markov fields. *Journal of Applied Probability*, 35(3):608–621, 1998.

[32] M J Symons. Clustering Criteria and Multivariate Normal Mixtures. *Biometrics*, 37(1):35, 1981.

[33] Geoffrey McLachlan and David Peel. ML Fitting of Mixture Models. In *Finite Mixture Models (Wiley Series in Probability and Statistics)*, Wiley Series in Probability and Statistics, pages 40–80. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2005.

[34] Douglas Steinley and Michael J Brusco. Evaluating mixture modeling for clustering: recommendations and cautions. *Psychological Methods*, 16(1):63–79, 2011.

[35] S. Kullback and R. a. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[36] Gilles Celeux and Gérard Govaert. Comparison of the mixture and the classification maximum likelihood in cluster analysis. *Journal of Statistical Computation and Simulation*, 47(3-4):127–146, 1993.

[37] G. Govaert and M. Nadif. Comparison of the mixture and the classification maximum likelihood in cluster analysis with binary data. *Computational Statistics and Data Analysis*, 23(1):65–81, 1996.

[38] Daniel J Bauer and Patrick J Curran. Distributional assumptions of growth mixture models: implications for overextraction of latent trajectory classes. *Psychological methods*, 8(3):338–363, 2003.

[39] Peter G. Bryant. Large-sample results for optimization-based clustering methods. *Journal of Classification*, 8(1):31–44, 1991.

[40] Douglas Steinley and Michael J. Brusco. K-means clustering and mixture model clustering: Reply to McLachlan (2011) and Vermunt (2011). *Psychological Methods*, 16(1):89–92, 2011.

[41] John R Hipp and Daniel J Bauer. Local solutions in the estimation of growth mixture models. *Psychological Methods*, 11(1):36–53, 2006.

[42] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of Classification*, 2(1):193–218, 1985.

[43] Douglas Steinley. Properties of the Hubert-Arabie adjusted Rand index. *Psychological methods*, 9(3):386–396, 2004.

[44] Mark S Handcock, David R Hunter, Carter T Butts, Steven M Goodreau, and Martina Morris. statnet: Software tools for the Statistical Modeling of Network Data. 2003.

[45] T M Achenbach. Manual for the Youth Self Report and 1991 Profile. *Burlington VT*, 1991.

[46] D. B. Rubin. Multiple Imputation for Nonresponse in Surveys. *Harvard University*, (JOHN WILEY & SONS):15–19, 1987.

[47] Christophe Biernacki, Gilles Celeux, and Gérard Govaert. An improvement of the NEC criterion for assessing the number of clusters in a mixture model. *Pattern Recognition Letters*, 20(3):267–272, 1999.

[48] G Celeux and G Soromenho. An entropy criterion for assessing the numbers of clusters in a mixture model. *J. Classif.*, 13:195–212, 1996.

[49] Jean-Patrick Baudry. Estimation and model selection for model-based clustering with the conditional classification likelihood. *Electronic Journal of Statistics*, 9(0):1041–1077, 2015.

[50] Sarah R Gilman, Ivan Iossifov, Dan Levy, Michael Ronemus, Michael Wigler, and Dennis Vitkup. Rare De Novo Variants Associated with Autism Implicate a Large Functional Network of Genes Involved in Formation and Function of Synapses. *Neuron*, 70(5):898–907, 2011.

[51] Tor D Wager, Jian Kang, Timothy D Johnson, Thomas E Nichols, Ajay B. Satpute, and Lisa Feldman Barrett. A Bayesian Model of Category-Specific Emotional Brain Responses. *PLOS Computational Biology*, 11(4):e1004066, 2015.